University of Pisa and Scuola Superiore Sant'Anna

Master Degree Program in Computer Science and Networking

MASTER THESIS PROPOSALS

2013

This document presents proposals for Master Theses in Computer Science and Networking.

Each thesis is organized in the context of one or more reference courses. Each proposal is presented in the section of the main reference course: if other reference courses are involved, the respective sections contain a link to the main one. For example, the Information Retrieval section contains a link to a joint thesis described in the Algorithm Engineering section.

The proposals are described in a synthetic form in order to give a first idea of the topics and approaches. In some cases a broad area is presented, in which several theses can be defined.

More detailed information can be asked directly to the supervisors and/or acquired using the indicated links.

Contents

1 Algorithm Engineering

1.1 Designing and engineering a Prefix-search data structure

Supervisor(s): Paolo Ferragina, Rossano Venturini

Reference course(s): Algorithm Engineering, Information Retrieval

Short description of objectives and contents:

The goal of the thesis is to design, implement and test a data structure for supporting prefix-searches over a dictionary of strings, as it occurs e.g. in auto-completion queries of modern search engines. The data structure should be designed to achieve I/O-efficiency and compression, in order to work on dictionaries of million/billion strings. The starting points are few recent results obtained by the supervisors which nonetheless need engineering and testing. For the testing, the student will be provided with few libraries that have been already developed by other researchers, in order to make a fair comparison of the achieved results.

2 Distributed Enabling Platforms

2.1 See Distributed Systems: Paradigms and Models, Thesis 4

2.2 See High Performance Computing, Thesis 1

3 Distributed Systems: Paradigms and Models

3.1 Implementation of a generic macro data flow engine in FastFlow

Supervisor: Marco Danelutto, Massimo Coppola

Reference courses: Distributed Systems: Paradigms and Models (SPM), Programming Tools for Parallel and Distributed Systems (SPD)

Short description of objectives and contents:

The main goal of the thesis is the implementation of a generic macro data flow engine suitable to a) represent b) transform and c) efficiently execute macro data flow graphs on multi/many core shared memory architectures. In particular:

- An high level representation of the macro data flow graphs must be designed, where macro data flow instructions are represented as function, input token list, output destination list triples with some kind of representation suitable to use any "function" written in the host programming language (C, C++) as the function to be computed.
- Suitable tools supporting macro data flow rewriting should be implemented, suitable to support the application programmer in the tuning of the application macro data flow code, and
- An efficient interpreter of macro data flow graphs has to be implemented, possibly exploiting all the interesting features of the underlying architecture (e.g. presence of co-processors (GPU, MIC), NUMA features, cache coherence mechanisms peculiarities).

Possible issues specifically developed in the reference courses:

In the SPM course, specific activities concerning the macro data flow model may be assigned to the student, partially fulfilling the requirements of the "course project". In the SPD course, aspects related to the exploitation of GPU and MIC coprocessors may be considered as integral part of the course project.

3.2 Implementation of a distributed structured parallel programming framework in Erlang

Supervisor: Marco Danelutto

Reference courses: Distributed Systems: Paradigms and Models (SPM)

Short description of objectives and contents:

The thesis aims at developing a structured parallel programming framework in Erlang, suitable to execute parallel Erlang applications on networked multicore Linux workstations. The implementation of the structured parallel programming framework will leverage the primitive and efficient mechanisms provided by Erlang to implement processes and the related message passing primitives as well as to implement cooperating instances of the Erlang interpreter on internetworked workstations. The thesis will start from preliminary experiments, prototypes and achievements from the ParaPhrase EU FP7 STREP project. The validation of the framework will be performed using workstations currently at the Dept. of

Computer Science including Infiniband interconnected Xeon multicore, GPU (Tesla, Kepler) equipped workstations and Xeon PHI equipped workstations.

Possible issues specifically developed in the reference courses:

Within the SPM course, the possibility to develop a project in Erlang rather than using the frameworks usually taken into account for the course will be considered, such that preliminary knowledge of Erlang could be acquired by the student during the course project.

3.3 Implementation of high level parallel design patterns in TBB and OpenMP

Supervisor: Marco Danelutto, Massimo Coppola

Reference courses: Distributed Systems: Paradigms and Models (SPM), Programming Tools for Parallel and Distributed Systems (SPD)

Short description of objectives and contents:

The main goal of the thesis is the development of a library of high level parallel design patterns on top of Intel TBB and OpenMP, implemented as algorithmic skeleton, that is ready to use object classes implementing the patterns on top of multi/many core architectures. Both TBB and OpenMP *de facto* already provide some specific patterns (e.g. the pipeline pattern in TBB or "kind of" map pattern (parallel for) in OpenMP). The idea here is to provide an homogeneous, general set of skeletons implementing the nonexistent patterns or wrapping the existent ones in uniform pattern class, with the final goal of providing the application programmer with a consistent and efficient set of patterns to implement applications on top of a wide range of different shared memory multi core architectures. As a secondary goal, the thesis will investigate the possibility to run part of the application skeletons (or part of a single skeleton) on MIC style coprocessors (e.g. Xeon PHI) in case one or more of these devices are present.

Possible issues specifically developed in the reference courses:

SPD course will consider to assign specific tasks relative to Xeon PHI programming with TBB/OpenMP to the student, in such a way the preliminary background needed to exploit this architecture is acquired during the project course.

3.4 A framework supporting the execution of FastFlow programs on cloud

Supervisor: Marco Danelutto, Nicola Tonellotto

Reference courses: Distributed Systems: Paradigms and Models (SPM), Distributed Enabling Platforms (PAD)

Short description of objectives and contents:

The main goal of the thesis is the development of a framework supporting the execution of FastFlow structured programs on a set of multicore cloud computing nodes, possibly equipped with GPUs. The thesis will build on already existing prototypes demonstrating the feasibility of the deployment of FastFlow virtual machines on cloud nodes implementing workers of a centrally managed task farm. The thesis will manage to implement:

- A generalized mechanism for the deployment of FastFlow virtual machines on the cloud
- Suitable tools to support the automatic creation of virtual machine images from the original, non cloud specific, FastFlow program

- Suitable tools to reconfigure the set of cloud deployed virtual machines if needed (e.g. adding more resources or getting rid of part of the already recruited resources, in case of load increase/decrease in the application or in case of under/over performing of the assigned resources).

The thesis will be assigned in the framework of the activities related to the FP7 EU STREP project ParaPhrase.

Possible issues specifically developed in the reference courses:

PAD activities/project assigned to the student will include topics such that the student may acquire the preliminary knowledge needed to tackle with the thesis goals. In particular, specific activities related to the automated deployment of an orchestrated set of virtual machines on the cloud resources will be investigated. SPM course will provide the FastFlow background necessary to deal with the thesis structured parallel programming topics.

4 High Performance Computing

4.1 Parallel Data Stream Processing

Supervisor(s): Marco Vanneschi, Nicola Tonellotto

Reference course(s): High-Performance Computing (HPC), Distribuited Enabling Platforms (PAD)

Short description of objectives and contents:

Data Stream Processing (DaSP) is an emerging application paradigm in which computations are applied to data that are dynamically and continuously flowing on streams, instead of being stored in permanent data structures. While in the basic stream processing computations each stream element is computed independently of the others, in the general case of DaSP more than one (possibly all, or subsets of) stream element(s) is needed at each step of the computation. A typical example is the join function applied to two streams, A and B: the function is applied to all the possible combinations of elements of A and B or to their proper subsets ("windows"). DaSP applications exist in many fields, such as finance, network control, disaster prevention, data mining, social networks. In a DaSP application data must be processed "on the fly", thus in real time, as they flow through the systems. Often there are serious performance problems (service time, response time) that limit the applicability of DaSP computations or that require a very (too) large amount of temporary memory. For this reason high-performance implementations are needed: while parallel paradigms for basic stream processing are largely known (see the HPC course), the Parallel DasP approach is much more complex and is a substantially open research problem.

The research group has started a research line on Parallel DaSP aiming to define and to implement proper parallelism paradigms and programming tools. In this thesis, the first results are applied to investigate one or more applications in order to gain information and knowledge about future DaSP programming models. Applications will be selected in the fields mentioned before, and will be experimented and evaluated on parallel machines (clusters with multicore nodes) available at the Laboratory of Parallel Architectures (Department of Computer Science).

A notable application is described in the joint thesis 5: see *Thesis 2 of the coordinating course Packet Switching and Processing Architectures*.

4.2 Adaptation strategies for Cloud applications

Supervisor(s): Marco Vanneschi, Massimo Coppola

Reference course(s): High-Performance Computing (HPC), Programming tools for Parallel and

Distributed Systems (SPD)

Short description of objectives and contents:

Autonomic High-Performance computing represents an emerging paradigm to develop distributed parallel applications in execution environments featuring a high variability of parameters such as calculation times of sequential functions, communication latencies as well as the Quality of Service (QoS) required by the users. In this scenario it is of great importance to adapt the application configuration in response to events that were unexpected when the application was firstly deployed on the target architecture. The adaptation process consists in developing parallel components capable of changing their configuration expressed in terms of

parallelism degree and parallelism form (structured parallelism patterns). The thesis is a part of an ongoing research where advanced formal control-theoretic strategies are being developed in order to optimize the adaptation process by: *i)* reaching desired QoS levels in face of variable execution conditions; *ii)* improving the stability/durability of control decisions; *iii)* reducing operating costs related to the application execution.

The approach adopted in the thesis is based on Model-based Predictive Control, a technique proposed by Gabriele Mencagli (PhD Thesis) and currently developed in our research group. The approach deals with reconfigurations by exploiting performance models of structured parallelism patterns and statistical predictions of disturbances over short future horizons.

The thesis inherits all the methodological background studied so far by our research group and is aimed at providing practical validations of the approach in real-world applications developed in Cloud environments. The candidate will participate to a synergic work between the Parallel Architecture Research Group at Department of Computer Science for the methodological part and the High-Performance Laboratory at ISTI CNR for the part of Cloud applications.

4.3 New structured parallel programming supports for manycore architectures

Supervisor(s): Marco Vanneschi

Reference course(s): High-Performance Computing (HPC)

Short description of objectives and contents:

Our research group in High Performance Computing and Programming Models is working on architectural and cost model for multicore-based parallel architectures, aiming at experimenting and formalizing new approaches able to match the parallel architecture with the structured parallel programming paradigms for high performance applications (see HPC course). Many research problems are related to the memory hierarchy organization of multicore based architectures, notably exploitation of locality and reuse, and communication overhead minimization, in relation with the properties of parallel computation having a graph structure composed of structured parallel modules.

The proposed theses in this area are studied from the methodological and from the experimental point of view, in order to derive new architectural and cost models, and new run-time support models of structured parallel paradigms. Rapid prototyping implementations, application/benchmarking experiments, and performance evaluations are done on the most recent manycore architectures in the Parallel Architecture Research Group at Department of Computer Science, specifically Tilera and Intel PHI.

Specific theses in this area are:

1) *Structured parallel programming paradigms on SMP- vs NUMA-like abstract architectures,*
2) *Structured parallel programming paradigms on manycore-manychip parallel architectures,*
3) *Structured parallel programming paradigms using automatic and non-automatic cache coherence.*

They will be based on the PhD Thesis work by Daniele Buono and Silvia Lametti, in both cases with the collaboration of Gabriele Mencagli (Post Doc) and Tiziano De Matteis (PhD student).

4.4 See Packet Switching and Processing Architectures, Thesis 1

4.5 See Packet Switching and Processing Architectures, Thesis 2

5 Information Retrieval

5.1 See Algorithm Engineering, Thesis 1

6 Network Optimization Methods

6.1 Delay-Constrained Routing Problems

Supervisor(s): Maria Grazia Scutellà, Antonio Frangioni

Reference course(s): Network Optimization Methods (MOR)

Short description of objectives and contents:

Real-time traffic with stringent Quality of Service requirements is becoming more and more prevalent in contemporary telecommunication networks. When maximum packet delay has to be considered, optimal delay-constrained routing requires not only choosing a path, but also reserving resources (transmission capacity) along its arcs, as the delay is a nonlinear function of both kinds of components. The thesis is a part of an ongoing project with computer network experts where innovative models of the problem are being developed together with efficient algorithms for their solution. The currently available models are formulated as a mixed-integer Second-Order Cone (SOCP) programs, and therefore can be solved with off-the-shelf technology; a particular class of reformulation techniques, known as Perspective Reformulation, can be used to tighten the continuous relaxation. Algorithm-wise, specialized algorithms have only been developed for simple versions of the problem where only one flow is routed, only one of the possible (nonlinear) delay formulae is considered, and all arcs are reserved the same capacity (this is referred to as ERA, i.e., Equal Rate Allocations). The thesis aims at developing algorithms for more complex versions of the problem, hopefully in the more complex multi-flow setting where the routing of all origin-destination pairs is decided simultaneously. There are several possible lines of research (approximation algorithms, Lagrangean relaxation and Branch-and-Price, combinatorial approaches such as dynamic programming), and therefore the candidate will have a choice about which of them to pursue; actually, more than one candidates could be accommodated, and some degrees of synergy are possible if the candidates are willing to collaborate.

Issues specifically developed in the reference courses (subject, course, approximate number of hours):

Delay-Constrained Routing, Second-Order Cone Programs, Convexification techniques, Approximation algorithms. MOR, 10 hours.

6.2 Experimental Evaluation of Joint Chance Constrained Formulations of Telecommunication Network Design Problems under Uncertainty

Supervisor(s): Maria Grazia Scutellà, Antonio Frangioni

Reference course(s): Network Optimization Methods (MOR)

Short description of objectives and contents:

When designing or upgrading a communication network, operators are faced with a major issue, as uncertainty on communication demands makes it difficult to correctly provision the network capacity. In general, what an operator has available (at best) is a collection of past realization of traffic matrices, that must be used to somehow predict the future traffic patterns; in turn, this prediction must be used to correctly provision the network. A conceptually simple approach is to extrapolate from past data some (approximate) probabilistic distribution of traffic matrices and then define a Chance Constrained Problem whereby one aims at finding the cheapest capacity allocation that guarantees, within a prescribed level of confidence, that

each arc will support the traffic demands. However, this turns out to be, in general, a difficult non convex optimization problem. In the individual case (i.e., when bounding the failure probability of each arc separately), it is possible to approximate the chance constraints with convex sets, leading to manageable optimization problems. This has been shown to be effective in practice. However, in principle one would rather control the joint probability failure across all arcs: this leads to more complex problems, but some recent developments (in particular decomposition algorithms belonging to the class of inexact bundle methods) may allow to solve real-world instances efficiently enough. The aim of the thesis is to study and implement a decomposition algorithm for the Joint Chance Constrained version of the problem, building on a good set of available software tools, and compare its effectiveness to the existing well-tested framework of Individual Chance Constrained codes, on data of real-world telecommunication networks.

Issues specifically developed in the reference courses (subject, course, approximate number of hours):

Second-Order Cone Programs, Chance Constrained Optimization, Decomposition algorithms. MOR, 10 hours.

6.3 Combinatorial Multigrid Approaches for Interior-Point Approaches to Min-Cost Flow Problems

Supervisor(s): Antonio Frangioni

Reference course(s): Network Optimization Methods (MOR)

Short description of objectives and contents:

The Min-Cost Flow Problem is a fundamental optimization problem with countless applications (often as a subproblem), in particular in the design and management of telecommunication networks. Among the possible approaches to the problem, Interior-Point ones are particularly attractive for very-large scale instances; however, they require at each step the solution of very large, graph-structured linear systems (KKT systems), that need to be tackled with sophisticated ad-hoc methods in order for the overall algorithm to be competitive. Actually, the matrices characterizing these systems belong to a larger class (that of L-matrices) that have numerous applications outside optimization (Web searching engines, Markov chains, consensus algorithms, ...) and therefore are extensively researched. After a first slew of results focussed on Preconditioned Conjugate Gradient methods, in recent years "combinatorial" Multigrid approaches have been extensively investigated, and very promising theoretical results have been obtained. However, these results have only very partly been computationally verified (against the best-of-breed of the previous approaches), especially in the optimization setting that has its particular quirks (repeated solution of the systems with the same graph structure but changing arc weights). The aim of the thesis is to review the recent literature on multigrid approaches for L-matrices and implement one or more of the most promising proposals, possibly exploiting available existing code (if any), to be tested in the context of an existing C++ software framework for IP approaches to structured Linear Programs (such as MCF).

Issues specifically developed in the reference courses (subject, course, approximate number of hours):

Interior-Point algorithms, linear algebra for graph-structured matrices. MOR, 10 hours.

6.4 Vectorization of Interior-Point Approaches to (Quadratic) Min-Cost Flow Problems

Supervisor(s): Antonio Frangioni

Reference course(s): Network Optimization Methods (MOR)

Short description of objectives and contents:

The Min-Cost Flow Problem is a fundamental optimization problem with countless applications (often as a subproblem), in particular in the design and management of telecommunication networks. Among the possible approaches to the problem, Interior-Point ones are attractive for very-large-scale instances; in particular, unlike all known "combinatorial" approaches they have the nice feature to be easily extendable to solve problems with a convex Quadratic separable objective function at basically no extra cost. The crucial step in these algorithms is the solution of very large graph-structured linear systems (KKT systems), that need be tackled with sophisticated ad-hoc methods in order for the overall algorithm to be competitive; however, a number of approaches have already been devised, for instance based on Preconditioned Conjugate Gradient methods using "chordal graph preconditiones" (mostly trees) or, more recently, "combinatorial" Multigrid methods. Almost all the costly steps of these algorithms involve simple vector operations; this means that a further benefit of the IP approach is to be, for the most part, easily "vectorized" when appropriate vector processing hardware (e.g. a GPU) is available. This is with the possible exception of a few "combinatorial" parts, such as the computation of the preconditioner (which basically requires the computation of maximal spanning trees); yet, these could also probably be vectorized, albeit not in a trivial way. The main aim of the thesis is to vectorize an existing C++ software framework for IP approaches to structured Linear Programs (such as MCF) using the appropriate techniques (CUDA or OpenCL); this is the "easy" part because all the involved operations are "naturally vectorizable". Once this is done, ideas will have to be proposed for the vectorization of the combinatorial part as well, in order to obtain a fully vectorized code that can be tested on large-scale (Q)MCF instances. This thesis could be peformed by two candidates, with one of the two taking on the more "technical" approach of vectorizing the code and the other researching on the more "methodological" issue of finding algorithms for maximal spanning trees that can be efficiently vectorized.

Issues specifically developed in the reference courses (subject, course, approximate number of hours):

Interior-Point algorithms, vectorization of graph algorithms. MOR, 10 hours.

6.5 Primal-Dual Approaches to Quadratic Min-Cost Flow Problems

Supervisor(s): Antonio Frangioni

Reference course(s): Network Optimization Methods (MOR)

Short description of objectives and contents:

The Min-Cost Flow Problem is a fundamental optimization problem with countless applications (often as a subproblem), in particular in the design and management of telecommunication networks. Several applications would require the solution of the (convex, separable) Quadratic version of the problem, but while many efficient combinatorial approaches for the standard linear case have been proposed (and well-engineered implementations of these are actually available), there are very few specialized algorithms for Quadratic MCF, and basically no actual solver available. Among the combinatorial algorithms

for (linear) MCF, primal-dual ones are perhaps the most efficient. The main aim of the thesis is to theoretically extend primal-dual approaches for the linear MCF to the Quadratic case, which is surely doable, although the theoretical and practical efficiency of these is difficult to predict at this point. Once this is done, a practical implementation should be developed which allows to at least verify the correctness of the algorithm and obtain insights on its computational behavior in practice; this should be relatively simple, at least for a prototypical implementation, by exploiting one or more of the relevant available C++ projects (MCFClass, LEMON, GTL). Depending on the results and on the available time, some further work could be dedicated to either theoretically improve the algorithm once its practical pitfalls are clarified, or to produce better engineered implementations that can compete with state-of-the-art solvers.

Issues specifically developed in the reference courses (subject, course, approximate number of hours):

Primal-dual algorithms for Min-Cost Flow Problems, implementation of graph algorithms in C++. MOR, 10 hours.

7 Network Security

7.1 Simulating Intelligent Attacks against an ICT Infrastructure

Supervisor(s): Fabrizio Baiardi

Reference course(s): Network Security (SR)

Rather than a single issue, the following proposal describes a rather broad area where each student, or groups of student, can select some specific topics to be investigated.

The broad area concerns the development of one or more tools to run a Monte Carlo simulation of complex attacks by intelligent agents against a complex ICT infrastructure such as a cloud system or a critical ICT infrastructure such as those in a SCADA control plan. Being intelligent, an agent can execute not only elementary attacks but also sequences of elementary attacks that implement a privilege escalation till the agent controls some resources of the target infrastructure. The goal of the simulation is to produce data to built statistics to investigate in a quantitative way the security of the considered infrastructure and to a quantitative methodology and tools to support the assessment and the management of the risk posed by ICT infrastructures.

In practice, the input of the simulation is a database with information about

a) the infrastructure components,

b) the potential and effective vulnerabilities of the components

c) the elementary attacks enabled by the vulnerabilities

d) the threat agents, the resource they can access and their goals, eg the components they aim to control.

Starting from this database, the tools simulates the behaviours of the agents in terms of how they compose the various elementary attacks into an attack plan to achieve their goals. By repeating several time the simulation, we can build a database with statistical samples that can be used to build statistics to evaluate several properties of interest such as the average time to attack the infrastructure, the average impact due to one or several agents or how to increase the overall robustness by removing one or more vulnerabilities.

Within this broad area there are several topics that may be investigated. Among those that concerns the security fields and that may be investigated in a graduation thesis, I mention, as an example:

a) which are the best strategies that an agent may adopt to plan the attacks to be composed or, in other words, which is the best strategy an agent can adopt to run a penetration test against an ICT infrastructure,

b) how to map the output of a vulnerability scanner into the input database of the simulation,

c) how to determine the vulnerabilities to be patched to minimize the overall impact due to the agents,

d) how to discover how an agent changes the strategy to compose elementary attacks in to plans as some vulnerabilities are patched to stop some elementary attacks.

A thesis in this area enables a student the best strategies to attack or defend a system, to discover the weakness of an ICT infrastructure or to evaluate the overall infrastructure security.

There is second class of theses that fairly independent of the security field and are focused on the development of a highly parallel version of the simulator to minimize the execution time so that we can simulate even large infrastructures. Our goal is the simulation of nation wide infrastructures including a number of nodes in the range 1000-10000. To appreciate the resulting complexity, consider that, in turn, each node can be decomposed into further components e.g. an OS, some hardware resources, some applications. Furthermore, also the components of the interconnection structure of the infrastructure have to be simulated. A sequential implementation of the simulator can be applied only to infrastructure modelled as at most a few hundreds components.

This second class of theses will use an highly parallel IBM system including 100 cores that is devoted to the project and has been granted by IBM in the framework of a Shared University Research after a world wide competition. A thesis in this class can be of interest to students that are interested in getting acquainted with the theoretical and practical problems posed by the simulation of a highly complex ICT infrastructure rather than in computer and network security topics . In this way a student may acquire a know how to be used in fields such computer system performance evaluations.

8 Networks and Technologies for Telecommunications

8.1 Control and Management architectures for Cloud Networking

Supervisor(s): Barbara Martini, Piero Castoldi

Reference course(s): Networks and Technologies for Telecommunications (RTT)

Rather than a single issue, the following proposal describes a rather broad area where each student, or groups of student, can select some specific topics to be investigated for their thesis.

The broad area of research concerns novel control and management architectures and paradigms for dynamic resource provisioning in Cloud Networking environment. Indeed, recent advances in software virtualization techniques (i.e., Hypervisor and Virtual Machine operation) and ever more popular cloud computing services are pushing towards convergent Cloud Data Center infrastructures acting as a unified pool of shared resources for many distributed services (e.g., remote storage, multimedia, high performance computing). In such a context, the network interconnection among and within Data Centers become crucial along with the role of telecom providers in the context of Data Centers infrastructures. Key to the whole process is a control and management system able to jointly manage network, computing and data storage elements while enabling dynamic network reconfiguration and QoS-guaranteed data transfer dictated by cloud service dynamics. Also, the emerging of the highly promising Software-Defined Network (SDN) paradigm offer new opportunities to telecom providers for a dynamic control of network data flows over heterogeneous systems while offering connectivity services with different granularities and dynamics. Finally, security concerns arise for possible confidentiality breaches due to the disclosure of information on service assets and business initiatives when service interactions and orders occurs among different providers, either service or telecom providers.

Within this broad research area the following topics could be investigated:

1) Virtualization-aware network resource provisioning in Cloud Data Center:

The oversubscription of DC interconnection links and the high volatility of Virtual Machine (VM) deployments require a flexible and agile control of the DC network infrastructure, even integrated with computing and storage resources. Different thesis are available for this topic, and specifically, on novel virtualization-aware networking functionality that enables VM placements in Data Centers considering the actual utilization of the DC interconnection links. To this purpose, SDN approach is considered to exploit highly rate of programmability and decide for the best VM placement decision based on flow-level statistics and prevent traffic performance degradation during VM traffic exchanges.

2) Inter-Data Center network service provisioning for on-demand Cloud Service delivery:

Reconfigurable and dynamically provisioned connectivity service set-up is required for flexible server-to-server communications across Cloud DCs. Different thesis are available in this topic, and specifically, on network management and control processes enabling the set-up of inter-DC connectivity services that is driven by cloud service dynamics and, thus, dynamically activated (and not statically pre-provisioned) as demanded by a service provider for the purpose of addressing specified DC resource optimizations, e.g., server

consolidation across DCs, or for enabling prompt cloud service recovery, e.g., fast workload migrations due to server faults.

3) Statistical analysis for confidentiality attack detection in Multi-Provider environments:

In a multi-provider environment, a resource provider might not pleased with possible breaches of salient information about its infrastructure capabilities, either network infrastructure or DC resource environment. Different thesis are available in this topic, and specifically, on technical solutions that guarantee the confidentiality by detecting apparently licit sequence of service requests (e.g., connectivity service request, DC resource cluster) while they are sly attempts to obtain confidential information of some provider capability (e.g., bandwidth capability for a telecom provider). The final goal is to detect malicious activities without affecting the normal operation of both service customers and providers using statistical methods based on anomaly detection process.

These research topics derives from research lines already carried out in the context of both EU-funded research projects and collaborations with industry. Both simulation and experimental thesis are available with the support of the Networks and Services Lab facilities and staff at TeCIP Institute. Based on obtained results, publications of papers reporting outcomes of thesis work are envisaged on International conference proceedings and journals.

8.2 Optical multicast: tree computation and spectrum assignment

Supervisor: Dr. Nicola Sambo

Reference course: Networks and Technologies for Telecommunications (RTT)

Short description of objectives and contents:

Emerging Internet applications for both business and entertainment purposes are continuously increasing, leading to a massive bandwidth demand. Several of these services (e.g., real-time high-definition IP-TV) may require the fast delivery of the same digital content to multiple destinations, i.e., multicast. Currently, multicast is supported at the IP layer. This implies electronic elaboration at each router that has to replicate the packets toward different output ports, thus consuming energy and increasing costs. The support of multicast directly in the optical domain is a target for reducing the amount of optical–electronic–optical conversions (thus, the network operational and capital expenditure) and energy consumption and for achieving effective network resource utilization.

This thesis will be built on node architectures supporting optical multicast, especially in the emerging flex-grid optical network infrastructures, which permit to achieve high spectrum efficiency. The use cases will be referred to emerging services such as real-time HD videos.

This thesis will have the following objectives: i) the proposal of effective algorithms for the computation of a multicast tree; ii) the proposal of spectrum assignment strategies for multicast connections; iii) the study of control plane implications to set up and tear down multicast connections in a flex-grid optical networks.

8.3 De-fragmentation techniques in flex-grid optical networks

Supervisor: Dr. Nicola Sambo

Reference course: Networks and Technologies for Telecommunications (RTT)

Short description of objectives and contents:

In flex-grid optical networks, especially in the presence of dynamic traffic, the spectrum of the links may result to be "fragmented", i.e. the available resources are not contiguous in the spectrum. Fragmentation may significantly affect the overall network efficiency, given that connections require contiguous portion of the spectrum to be accommodate (otherwise requests are blocked). Effective techniques for de-fragmentation (i.e., re-optimization) are then required to limit the wasting of spectrum resources. Recently, de-fragmentation techniques (such as push-pull) have been proposed, guaranteeing de-fragmentation without loss of data.

Research is still needed to identify the optical connections that have to be shifted in frequency or re-routed in order to re-organize the traffic with the aim of increasing the overall network throughput.

This thesis will start from the most promising de-fragmentation techniques at the physical layer and will study and propose novel de-fragmentation algorithms to increase network throughput. The objectives are the proposal of algorithms for: i) the selection of lightpaths to be re-routed or re-shifted in frequency; ii) the path computation algorithms for re-optimization; iii) the spectrum assignment for re-optimization. The study will be carried on in centralized and/or distributed network scenarios.

8.4 Software defined networking for optical networks

Supervisor: Dr. Nicola Sambo

Reference course: Networks and Technologies for Telecommunications (RTT)

Short description of objectives and contents:

OpenFlow, and software defined networking (SDN) in general, have been getting a lot of attention during the last years. OpenFlow is an open standard based on a centralized control model, which is built on the fact that the most modern switches and routers (thus, nodes in circuit-switched networks) can be abstracted by identifying a common set of functions. In this context, OpenFlow has been defined as a protocol to program and set flow tables and cross-connects in a centralized manner. SDN and OpenFlow are seen as an attractive way to control circuit-switched optical networks. However, several issues are still open, such as the reliability in the presence of OpenFlow control.

This thesis will have the objective of proposing possible extensions to the OpenFlow architecture for optical networks, with particular attention to the network reliability.

9 Packet-Switching and Processing Architectures

9.1 Basic Openflow Support for PFQ

Supervisor(s): Gregorio Procissi, Nicola Bonelli, Stefano Giordano, Marco Vanneschi

Reference course(s): Packet-Switching and Processing Architectures (AED), Teletraffic Engineering (IT), High Performance Computing (HPC)

Short description of objectives and contents:

PFQ is a monitoring platform developed at the NetGroup Lab. of the University of Pisa that runs on commodity PCs on top of vanilla drivers for any network device and takes full advantage of CPU parallelism and multi–queue network cards. Currently, PFQ proves to capture packets at line speed with Intel 10 Gbps cards. At the high level view, PFQ is a Linux kernel module that taps packets from one or more network cards and distribute them to user–space sockets according to applications requests. Packets retrieved from multiple devices traverse *Steering functions* that provide a first–stage for data pre–processing, classification and forwarding towards upstream applications. As a special case of data pre-processing, steering functions can be used to implement OpenFlow functionalities.

The aims of this thesis twofold:

1. Design and develop a user space interface towards the Openflow protocol
2. Design and develop a specific steering function (at kernel space) to implement some basic Openflow directives

Possible issues specifically developed in the reference courses

The thesis requires skills on C, C++ and Linux kernel programming. OpenFlow and Network monitoring techniques are presented in the AED course.

Notice: this thesis is actually dimensioned for <u>a joint work of two students</u> working on the previously listed points 1 and 2, respectively.

9.2 Randomized Packet Filtering for PFQ

Supervisor(s): Gregorio Procissi, Nicola Bonelli, Marco Vanneschi

Reference course(s): Packet-Switching and Processing Architectures (AED), High Performance Computing (HPC)

Short description of objectives and contents:

PFQ is a monitoring platform developed at the NetGroup Lab. of the University of Pisa that runs on commodity PCs on top of vanilla drivers for any network device and takes full advantage of CPU parallelism and multi–queue network cards. Currently, PFQ proves to capture packets at line speed with Intel 10 Gbps cards. At the high level view, PFQ is a Linux kernel module that taps packets from one or more network cards and distribute them to user–space sockets according to applications requests. Packets retrieved from multiple devices traverse *Steering functions* that provide a first–stage for data pre–processing, filtering, classification and forwarding towards upstream applications.

In high speed traffic processing, a key issue is to immediately detect potentially interesting packets. At very high speed, this operation is particularly crucial as filtering packets *close to the wire* relieves real applications from handling large volumes of (uninteresting) data.

This thesis aims at integrating on PFQ a fast and randomized approach to packet filtering based on partitioning database filtering rules for their storage in fast and compact Bloom filters.

Operatively the thesis consists of:

1. designing and developing a user space application that partitions the filtering ruleset and builds the corresponding Bloom filters;
2. developing the (kernel space) steering function that performs [packet filtering *on-the-fly*

Possible issues specifically developed in the reference courses

The thesis requires skills on C, C++ and Linux kernel programming. Packet filtering, classification and network monitoring techniques are presented within the AED course.

10 Peer to Peer Systems

10.1 Anti-cheating Techniques for P2P Massively Multi-player Games

Supervisor: Laura Ricci.

Reference Course: Peer to Peer Systems (P2P).

Short description of objectives and contents:

A common trend in the development of distributed architectures for massively multi-player on-line games (MMOGs) is to exploit direct communications between players (peers) to decrease the load on centralized servers. The peers build and maintain a decentralized overlay which is exploited to notify the game events. One of the main problems of this approach is to foresee and avoid cheating, i.e. incorrect behaviour of players not following the game rules to increase their self-advantages. The main goal of this thesis is to project, develop and test a set of anti-cheating techniques for P2P MMOGs. The techniques will be developed within an existing P2P network, based on a gossip-based overlay.

The student will take into account both classic concepts like information authenticity and integrity and concepts more related to the specific MMOGs applications as well.

10.2 Integration of P2P and Cloud Architectures for Massively Multiplayer Games

Supervisor: Laura Ricci

Reference Course: Peer to Peer Systems (P2P).

Short description of objectives and contents:

A new trend in the definition of distributed architectures for massively multi-player on-line games (MMOGs), is to exploit heterogeneous resources, like P2P and cloud resources. In this context, this thesis will extend an existing architecture based on the integration of centralized servers and P2P nodes with the goal to consider cloud resources, instead of classical centralized servers. In this context, the student will propose and implement mechanisms to guarantee the given level of Quality of Service (Qos) for the application, while limiting the economical cost for the game provider. The problem of maintaining a given level of QoS is particularly challenging for interactive applications like MMOGs.

The thesis will investigate a set load prediction mechanisms based on different mathematical techniques (time-series, smoothing,....).

10.3 Definition of P2POverlays based on Delaunay Triangulations

Supervisor: Laura Ricci

Reference Course: Peer To Peer Systems (P2P).

Short description of objectives and contents:

Several P2P applications are based on the definition of a mapping of the peers into a virtual space where each peer is characterized by its cartesian coordinates, for instance a P2P mobile application may exploit geo-referential coordinates of the nodes. Some recent works propose to define Delaunay-based P2P overlays for this kind of applications. Delaunay triangulations are mathematical structures often exploited in computational geometry

modelling the local neighbourhood of the sites in a cartesian space. P2P Delaunay-based overlays support efficient geometrical routing algorithms, like compass routing. The main goal of this thesis is to extend existing approaches taking into account the scenario where the positions of the pees dynamically change over time. For instance, this is the case of nodes in a mobile network or of the avatars moving within a virtual space.

The student will evaluate the possibility of defining an approximated Delaunay triangulation taking into account the inconsistencies in the local view of the peers which may arise because of the network dinamicity.

11 Programming Tools for Parallel and Distributed Systems

11.1 See Distributed Systems: Paradigms and Models, Thesis 1

11.2 See Distributed Systems: Paradigms and Models, Thesis 3

11.3 See High Performance Computing, Thesis 2

12 Software Service Engineering

12.1 Deploying application modules over multiple clouds

Supervisor(s): Antonio Brogi

Reference course(s): Software Service Engineering (ISS)

Short description of objectives and contents:

How to deploy and manage, in an efficient and adaptive way, complex multi-service applications over technologically dissimilar cloud environments is one of the problems that have emerged with the cloud revolution. The objective of the thesis is to study suitable techniques to determine a distribution of application modules onto multiple available clouds while respecting QoS (Quality of Service) properties and technology requirements needed for individual application modules.

12.2 SLA-based service composition

Reference course(s): Software Service Engineering

Short description of objectives and contents:

Service-Level Agreements (SLAs) are intended to formalize service usage in measurable terms, for instance by specifying QoS (Quality of Service) indicators. Clearly, the SLA of a service composition relates to the SLAs of the composed services. The objective of the thesis is to study suitable techniques to assist the (expensive and error-prone) task of specifying "coherent" SLAs for service compositions.

12.3 Service adaptation

Reference course(s): Software Service Engineering

Short description of objectives and contents:

Service adaptation remains one of the main issues for application integration in a variety of situations, ranging from the need of overcoming mismatches among services developed by different parties, of customising existing services to different types of clients, or of adapting legacy systems to meet new business demands. The objective of the master thesis is to prototype the deployment of adapters capable of dynamically overcoming signature and behaviour mismatches in dynamic and resource-bounded environments.

13 Teletraffic Engineering

13.1 See Packet Switching and Processing Architectures, Thesis 1

14 Wireless Networks of Embedded Systems

Supervisor(s): Paolo Pagano

Reference course(s): Wireless Networks of Embedded Systems (RWE)

Rather than single issues, some of the following proposals describe set of theses referring to specific areas.

All theses will be done within the research activities carried on by the Networks of Embedded Systems area at the TeCIP institute: http://rtn.sssup.it/.

For more information about the theses: http://rtn.sssup.it/index.php/theses/.

For more information on hardware development boards: http://rtn.sssup.it/index.php/hardware/seed-eye

14.1 Multimedia Sensor Networks

WSN are usually deployed through a set of embedded devices, equipped by constrained resources (like computing capabilities and resident memory) and interconnected by a low rate, unreliable wireless network. Multimedia sensing is a challenging perspective to give "Eyes and Ears" to sensor devices and create added-value to distributed applications in ad-hoc networks.

Specific theses on this topic are:

1) *Multimedia Sensor Networks: Design and Implementation of a Video Codec for Embedded Systems*
2) *Multimedia Sensor Networks: Design and implementation of compressive sampling techniques*
3) *Multimedia Sensor Networks: Software alignment techniques for smart camera networks*
4) *Multimedia Sensor Networks: a Bandwidth allocation protocol targeted to Speech Communications*
5) *Multimedia Sensor Networks: Porting of the GSM AMR speech coder to a microcontroller architecture*

14.2 Localization protocols: WSN and RFID hybrid solution for indoor localization

This thesis consists in the design and implementation of a WSN and RFID hybrid solution to support localization services in the indoor environment. Low-complexity RFID reader will be integrated in WSN nodes to act as anchor of the localization system, thus overcoming the limits of an RSSI-based approach. Moreover, inertial sensors and PIR will be used to estimate the position when system anchors are not available.

14.3 Simulation platforms: A simple computer vision distributed protocol in WiSE-MNet

The goal addressed by WiSE-MNet is to estimate the impact of a real network over a distributed algorithm through simulation. However the logic inside the node is not exploited: in this thesis we propose to integrate Virtual Machines (emulating the ERIKA and Contiki micro-kernels) into

WiSE-MNet. In this situation it will be possible to assess the performances at node and network levels, i.e. on-board and distributed capabilities.

To test this complete simulator we propose to implement a simple computer vision distributed application**.**

14.4 Internet of Things: Multicast protocol for 6LoWPANs

WSNs consist of low-cost autonomous sensor devices which interact with each other in a wireless distributed system. Each sensor has very limited battery capacity, limited processor capability and limited storage capacity. Multicast paradigm reduces the communication costs for applications that send the same data to multiple recipients: instead of sending via multiple unicasts, multicasting minimizes the link bandwidth consumption, sender and router processing, and delivery delay.

The thesis will design and implement a Multicast protocol for Wireless Sensor Network. The main goal is to define a native multicast support for 6LoWPAN network (IPv6 over Lossy and Low Power Network, IETF 4919 standard). The experimental validation activity will be performed in a real scenario.

14.5 Intelligent Transport Systems

Vehicular Ad-hoc NETworks (VANET) are a special instance of nomadic networks where the mobile entities (i.e. cars) are expected to exchange information (vehicle to vehicle -- V2V) via the wireless channel. VANET technologies are considered as the basis of Intelligent Transport Systems; yet another challenge is that of interconnecting vehicular equipment with the road-side network (vehicle to infrastructure - V2I).

Our Seed-Eye boards are being customized to implement the ITS stations functionality (as standardized by ISO and ETSI) in tiny devices (like WSN). This activity has been preliminarily approved for standardization purposes as ISO Work Items #19079 and #19080.

Although Radio Frequency is the preferred technology for enabling V2I and V2V communication (see for instance the M/453 mandate by the European Commission to ETSI and CEN), alternative communication means are being promoted by the scientific community. A notable example is offered by Visible Light Communication (VLC).

Specific theses on this topic are:

1) *Intelligent Transport Systems: IPv6 adaptation to IEEE802.15.7-compliant Visible-light Personal Area Networks*
2) *Intelligent Transport Systems: IEEE802.15.4 and IEEE802.15.7 multi-mac networking for ITS-Stations*

14.6 Abstraction of Wireless Sensor Networks

As applications become more and more interconnected and interdependent, the number of objects, users and devices tends to increase. This poses the problem of the scalability of the communication and object management algorithms, and increases the complexity of administration. Ubiquitous Computing is a vision of the near future, in which an increasing number of devices embedded in various physical objects will be participating in a global information network. To set up a common ground of abstraction, a middleware layer should hide the heterogeneity of the network and the complexity of services and applications.

A prominent objective is that of implementing a code execution service at the node level by implementing a virtual machine capable of executing scripts or bytecodes.

Our Seed-Eye boards implement the Pymite interpreter to support the execution of scripts.

Specific theses on this topic:

1) *WSN abstractions: Rapid prototyping of distributed systems in WSN*
2) *WSN abstractions: An instrumentation layer for the grids*